Broadcast Using Certified Propagation Algorithm in Presence of Byzantine Faults[1]

Lewis Tseng, Nitin Vaidya, Vartika Bhandari

Abstract

We explore the correctness of the Certified Propagation Algorithm (CPA) [6, 1, 8, 5] in solving broadcast with locally bounded Byzantine faults. CPA allows the nodes to use only local information regarding the network topology. We provide a *tight* necessary and sufficient condition on the network topology for the correctness of CPA.

Keywords:
Distributed computing, Byzantine broadcast, CPA, Tight condition

1. Introduction

In this work, we explore fault-tolerant broadcast with locally bounded Byzantine faults in synchronous point-to-point networks. We assume a *f-locally bounded model*, in which at most f Byzantine faults occur in the neighborhood of every *fault-free* node [6]. In particular, we are interested in the necessary and sufficient condition on the underlying communication network topology for the correctness of the Certified Propagation Algorithm

[1] This research is supported in part by Army Research Office grant W-911-NF-0710287. Any opinions, findings, and conclusions or recommendations expressed here are those of the authors and do not necessarily reflect the views of the funding agencies or the U.S. government.

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **23 NOV 2014**	2. REPORT TYPE	3. DATES COVERED **00-00-2014 to 00-00-2014**

4. TITLE AND SUBTITLE **Broadcast Using Certified Propagation Algorithm in Presence of Byzantine Faults**	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **University of Illinois at Urbana-Champaign ,Coordinated Science Laboratory,1308 West Main Street,Urbana,IL,61801**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR’S ACRONYM(S)
	11. SPONSOR/MONITOR’S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT

We explore the correctness of the Certified Propagation Algorithm (CPA) [6, 1, 8, 5] in solving broadcast with locally bounded Byzantine faults. CPA allows the nodes to use only local information regarding the network topology. We provide a tight necessary and sufficient condition on the network topology for the correctness of CPA.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**9**	

(CPA) – the CPA algorithm has been analyzed in prior work [6, 1, 8, 5, 7].

Problem Formulation. Consider an arbitrary *directed* network of n nodes. One node in the network, called the *source* (s), is given an initial input, which the source node needs to transmit to all the other nodes. The source s is assumed to be *fault-free*. We say that CPA is *correct*, if it satisfies the following properties, where x_s denotes the input at source node s:

- **Termination:** every fault-free node i eventually decides on an output value y_i.
- **Validity:** for every fault-free node i, its output value y_i equals the source's input, i.e., $y_i = x_s$.

We study the condition on the network topology for the correctness of CPA.

Related Work. Several researchers have addressed CPA problem. [6] studied the problem in an infinite grid. [1] developed a sufficient condition in the context of arbitrary network topologies, but the sufficient condition proposed is not tight. [8] provided necessary and sufficient conditions, but the two conditions are not identical (not tight). [5] provided another condition that can approximate (within a factor of 2) the largest f for which CPA is correct in a given graph. Independently, [7] presented the tight condition in *undirected* graphs. Similar condition under other contexts are also discovered by other researchers [9, 3]. Please refer to [11] for more discussions.

System Model. The synchronous communication network consisting of n nodes including source node s is modeled as a simple *directed* graph $G(\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is the set of n nodes, and $\mathcal{E}$ is the set of directed edges between the nodes in $\mathcal{V}$. Node i can transmit messages to another node j if and only if the directed edge (i, j) is in $\mathcal{E}$. Each node can transmit messages to itself as well; however, for convenience, we exclude self-loops from set $\mathcal{E}$. That is, $(i, i) \notin \mathcal{E}$ for $i \in \mathcal{V}$. All the links (i.e., communication channels) are assumed to be point-to-point, reliable, FIFO (first-in first-out) and deliver each transmitted message exactly once. With a slight abuse of terminology, we will use the terms *edge* and *link* interchangeably.

For each node i, let N_i^- be the set of nodes from which i has incoming edges, i.e., $N_i^- = \{\, j \mid (j, i) \in \mathcal{E} \,\}$. Similarly, define N_i^+ as the set of nodes to which node i has outgoing edges, i.e., $N_i^+ = \{\, j \mid (i, j) \in \mathcal{E} \,\}$. Nodes in N_i^- and N_i^+ are, respectively, said to be incoming and outgoing neighbors of node i. Since we exclude self-loops from $\mathcal{E}$, $i \notin N_i^-$ and $i \notin N_i^+$. However, we note again that each node can indeed transmit messages to itself.

We consider the f-local fault model, with at most f incoming neighbors of any fault-free node becoming faulty. [6, 1, 8, 5, 7] also explored this fault model. Yet, to the best of our knowledge, the tight necessary and sufficient conditions for the correctness of CPA in *directed* networks under f-local fault model have not been developed previously.

2. Feasibility of CPA under f-local fault model

Certified Propagation Algorithm (CPA). We first describe the Certified Propagation Algorithm (CPA) from [6] formally. Note that the faulty nodes may deviate from this specification arbitrarily. Possible misbehavior includes sending incorrect and mismatching messages to different outgoing neighbors.

Source node s commits to its input x_s at the start of the algorithm, i.e., sets its output equal to x_s. The source node is said to have committed to x_s in round 0. The algorithm for each round r ($r > 0$), is as follows:

1. Each node that commits in round $r-1$ to some value x, transmits message x to all its outgoing neighbors, and then terminates.
2. If any node receives message x directly from source s, it commits to output x.
3. Through round r, if a node has received messages containing value x from at least $f+1$ distinct incoming neighbors, then it commits to output x.

The Necessary Condition. For CPA to be correct, the network graph $G(\mathcal{V}, \mathcal{E})$ must satisfy the necessary condition proved in this section. We borrow two relations $\Rightarrow$ and $\not\Rightarrow$ from our previous paper [12].

Definition 1. *For non-empty disjoint sets of nodes A and B,*

- *$A \Rightarrow B$ iff there exists a node $v \in B$ that has at least $f+1$ distinct incoming neighbors in A, i.e., $|N_v^- \cap A| > f$.*
- *$A \not\Rightarrow B$ iff $A \Rightarrow B$ is not true.*

Definition 2. *Set $F \subseteq \mathcal{V}$ is said to be a feasible f-local fault set, if for each node $v \notin F$, F contains at most f incoming neighbors of node v. That is, for every $v \in \mathcal{V} - F, |N_v^- \cap F| \leq f$.*

We now derive the necessary condition on the network topology.

Theorem 1. *Suppose that CPA is correct in graph $G(\mathcal{V}, \mathcal{E})$ under the f-local fault model. Let sets F, L, R form a partition[2] of $\mathcal{V}$, such that (i) source $s \in L$, (ii) R is non-empty, and (iii) F is a feasible f-local fault set. Then*

- *$L \Rightarrow R$, or*
- *R contains an outgoing neighbor of s, i.e., $N_s^+ \cap R \neq \emptyset$.*

Proof. The proof is by contradiction. Consider any partition F, L, R such that $s \in L$, R is non-empty, and F is a feasible f-local fault set. Suppose that the input at s is x_s. Consider any single execution of the CPA algorithm such that the nodes in F behave as if they have crashed.

By assumption, CPA is correct in the given network under such a behavior by the faulty nodes. Thus, all the fault-free nodes eventually commit their output to x_s. Let round r ($r > 0$), be the earliest round in which at least one of the nodes in R commits to x_s. Let v be one of the node in R that commits in round r. Such a node v must exist since R is non-empty, and it does not contain source node s. For node v to be able to commit, as per specification of the CPA algorithm, either node v should receive the message x_s directly from the source s, or node v must have $f + 1$ distinct incoming neighbors that have already committed to x_s. By definition of node v, nodes that have committed to x_s prior to v must be outside R; since nodes in F behave as crashed, these $f + 1$ nodes must be in L. Thus, either $(s, v) \in \mathcal{E}$, or node v has at least $f + 1$ distinct incoming neighbors in set L.

□

[2] Sets $X_1, X_2, X_3, ..., X_p$ are said to form a partition of set X provided that (i) $\cup_{1 \leq i \leq p} X_i = X$, and (ii) $X_i \cap X_j = \Phi$ if $i \neq j$.

Sufficiency. We now show that the condition in Theorem 1 is also sufficient.

Theorem 2. *If $G(\mathcal{V}, \mathcal{E})$ satisfies the condition in Theorem 1, then CPA is correct in $G(\mathcal{V}, \mathcal{E})$ under the f-local fault model.*

Proof. Suppose that $G(\mathcal{V}, \mathcal{E})$ satisfies the condition in Theorem 1. Let F' be the set of faulty nodes. By assumption, F' is a feasible local fault set. Let x_s be the input at source node s. We will show that, (i) fault-free nodes do not commit to any value other than x_s (Validity), and, (ii) until all the fault-free nodes have committed, in each round of CPA, at least one additional fault-free node commits to value x_s (Termination). The proof is by induction.
Induction basis: Source node s commits in round 0 to output equal to its input x_s. No other fault-free nodes commit in round 0.
Induction: Suppose that L is the set of fault-free nodes that have committed to x_s through round r, $r \geq 0$. Thus, $s \in L$. Define $R = \mathcal{V} - L - F'$. If $R = \emptyset$, then the proof is complete. Let us now assume that $R \neq \emptyset$.

Now consider round $r + 1$.

- Validity:

  Consider any fault-free node u that has not committed prior to round $r+1$ (i.e., $u \in R$). All the nodes in L have committed to x_s by the end of round r. Thus, in round $r+1$ or earlier, node u may receive messages containing values different from x_s only from nodes in F'. Since there are at most f incoming neighbors of u in F', node u cannot commit to any value different from x_s in round $r + 1$.

- Termination:

By the condition in Theorem 1, there exists a node w in R such that (i) node w has an incoming link from s, or (ii) node w has incoming links from $f+1$ nodes in L. In case (i), node w will commit to x_s on receiving x_s from node s in round $r+1$ (in fact, $r+1$ in this case must be 1). In case (ii), first observe that all the nodes in L from whom node w has incoming links have committed to x_s (by definition of L). Then, node w will be able to commit to x_s after receiving messages from at least $f+1$ incoming neighbors in L, since all nodes in L have committed to x_s by the end of round r by the definition of L.[3] Thus, node w will commit to x_s in round $r+1$.

This completes the proof. □

3. Discussion

This section presents extensions and complexity of verifying the condition. Due to space limitation, please refer to [11] for details.

CPA without prior knowledge of f. Given a graph G that can tolerate f-local faults (where f is unknown), we construct a broadcast algorithm in G without usage of f. The core idea is for each node to exhaustively test all possible parameters by running $n+1$ instances of CPA algorithm in parallel.

Other Communication Model. In the broadcast model [6, 1], when a node transmits a value, all of its outgoing neighbors receive this value identically.

[3] Since node w did not commit prior to round $r+1$, it follows that at least one node in L must have committed in round r.

Thus, no node can transmit mismatching values to different outgoing neighbors. In the asynchronous model [2], the algorithm may not proceed in rounds, but a node still commits to value x either on receiving the value directly from s, or from $f+1$ nodes. Under both models, condition in Theorem 1 is both necessary and sufficient for the correctness of CPA. The claim for asynchronous model may seem to contradict the FLP result [4]. However, our claim assumes that the source node is fault-free, unlike [4].

Complexity. [7] proved that it is NP-hard to examine whether CPA is correct in a given *undirected* graph. The condition in [7] is indeed equivalent to our condition (condition in Theorem 1) in *undirected* graphs. Therefore, it is NP-hard to examine whether a given graph satisfies our condition or not.

4. Conclusion

In this paper, we explore broadcast in arbitrary network using the CPA algorithm in f-local fault model. In particular, we provide a *tight* necessary and sufficient condition on the underlying network for the correctness of CPA.

References

[1] V. Bhandari and N. H. Vaidya. On reliable broadcast in a radio network: A simplified characterization. Technical report, UIUC, 2005.

[2] D. Dolev, N. Lynch, S. Pinter, E. Stark, and W. Weihl. Reaching approximate agreement in the presence of faults. *J. ACM*, 1986.

[3] D. Easley and J. Kleinberg. Networks, crowds, and markets: reasoning about a highly connected world. Cambridge, 2010.

[4] M. Fischer, N. Lynch, and M. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 1985.

[5] A. Ichimura and M. Shigeno. A new parameter for a broadcast algorithm with locally bounded Byzantine faults. IPL, 2010.

[6] C.-Y. Koo. Broadcast in radio networks tolerating Byzantine adversarial behavior. PODC, 2004.

[7] A. Pagourtzis, G. Panagiotakos, and D. Sakavalas. Reliable broadcast with respect to topology knowledge. DISC, 2014.

[8] A. Pelc and D. Peleg. Broadcasting with locally bounded Byzantine faults. IPL, 2005.

[9] N. B. Shah, K. V. Rashmi, and K. Ramchandran. Efficient and distributed secret sharing in general networks. CoRR abs/1207.0120, 2012.

[10] L. Tseng and N. Vaidya. Iterative approximate byzantine consensus under a generalized fault model. ICDCN, 2013.

[11] L. Tseng, N.Vaidya, V. Bhandari Broadcast using certified propagation algorithm in presence of Byzantine faults. CoRR abs/1209.4620, 2013.

[12] N. Vaidya, L. Tseng, and G. Liang. Iterative approximate Byzantine consensus in arbitrary directed graphs. PODC, 2012.